



# Qualcomm Centriq™ 2400 Processor

---

**Barry Wolford, Senior Director, Engineering**  
**Thomas Speier, Senior Director, Engineering**  
**Dileep Bhandarkar, Vice President, Technology**

Qualcomm Datacenter Technologies, Inc.

August 22, 2017

@qualcomm

Qualcomm Centriq 2400 Processor is a product of Qualcomm Datacenter Technologies, Inc.

# Agenda

---

- Qualcomm Datacenter Technologies Introduction
- Qualcomm® Falkor™ CPU Overview
- Qualcomm Centriq™ 2400 Server SoC Overview
- Summary

# QDT Well Positioned to Address Cloud Datacenter Opportunity

## Unique High Performance, Low Power ARM Based CPUs



Bringing decade of experience delivering high-performance, power-efficient ARM CPU architectures




Focus on true server class features and performance with aggressive power management techniques



Partnering with cloud market leaders for product definition



Uniquely positioned to leverage process leadership driven by mobile industry growth to deliver industry first 10 nm server processor



# Qualcomm Falkor™ CPU Designed for the Cloud

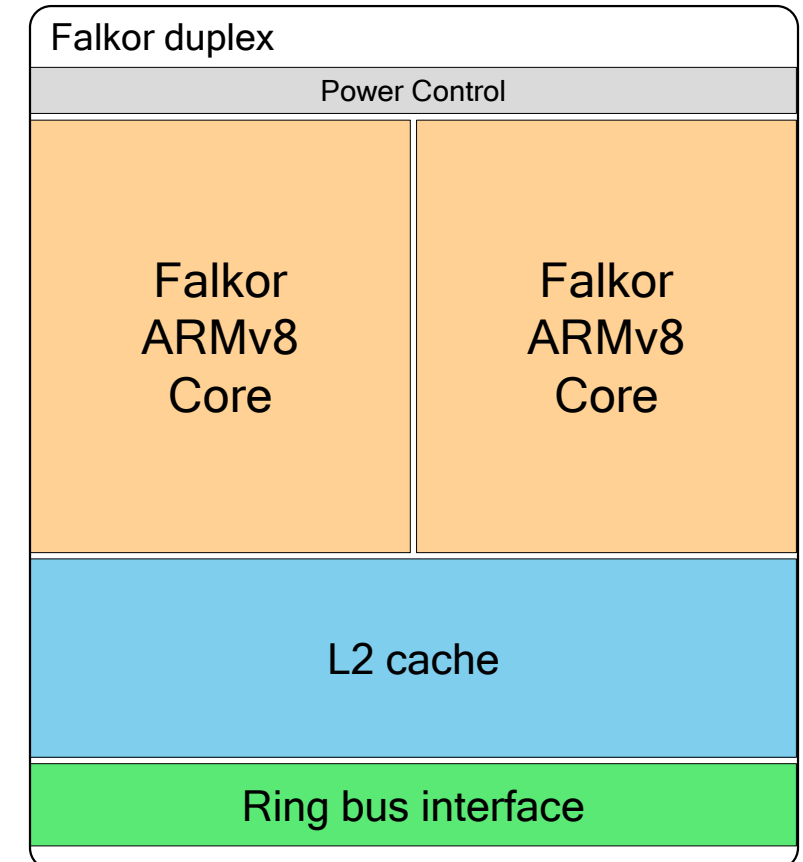
---

- QDT-designed custom core powering Qualcomm Centriq 2400 Processor
- 5<sup>th</sup> generation custom core design
- Designed from the ground up to meet the needs of cloud service providers
- Fully ARMv8-compliant
- AArch64 only
- Supports EL3 (TrustZone) and EL2 (hypervisor)
- Includes optional cryptography acceleration instructions
  - AES, SHA1, SHA2-256
- Designed for performance, optimized for power




# Falkor Core configuration

- Falkor core duplex is building block for SoC
- Two Custom ARM V8 CPUs
- Shared L2 Cache
- Nominal Operating Voltage ~1V
- Shared bus interface to Qualcomm® System Bus (QSB) ring interconnect
  - ❑ Qualcomm Proprietary Protocol
  - ❑ Custom Bi-Directional Segmented Ring Bus
    - ❑ Fully Coherent (Cache & IO)
    - ❑ Shortest Path Routing
    - ❑ Multicast on Read
    - ❑ > 250 GB/s aggregate bandwidth



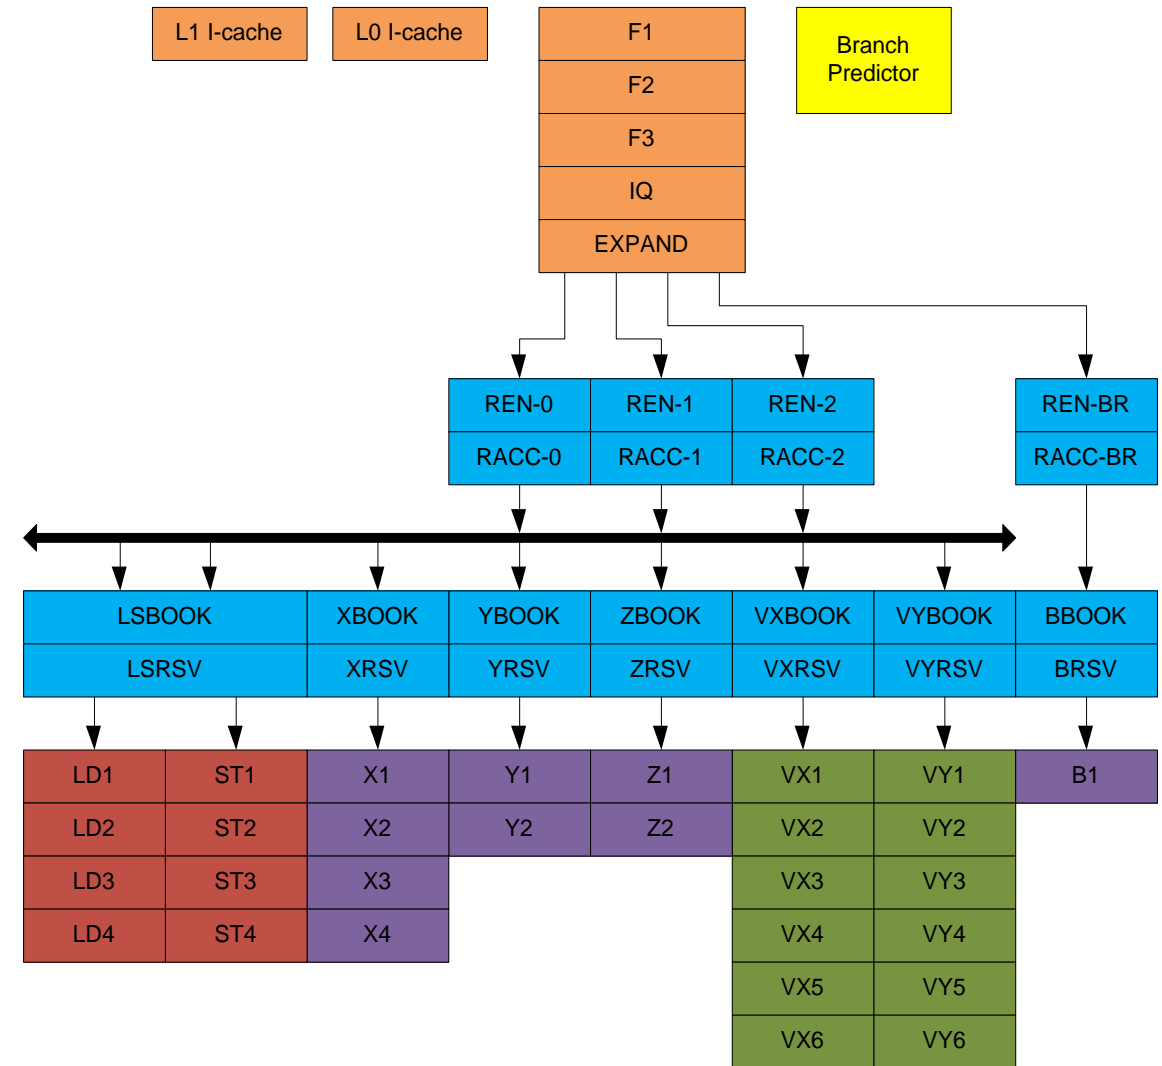

# Falkor L2 Cache

- 128-byte lines, 8-way
- Unified between I-side and D-side
- Shared between two CPUs in duplex
- 128-byte interleaved for improved throughput
- SEC-DED ECC protected
- 15-cycle minimum latency for L2 hit
- Inclusive of L1 D-caches
- 32-bytes per direction per interleave per cycle



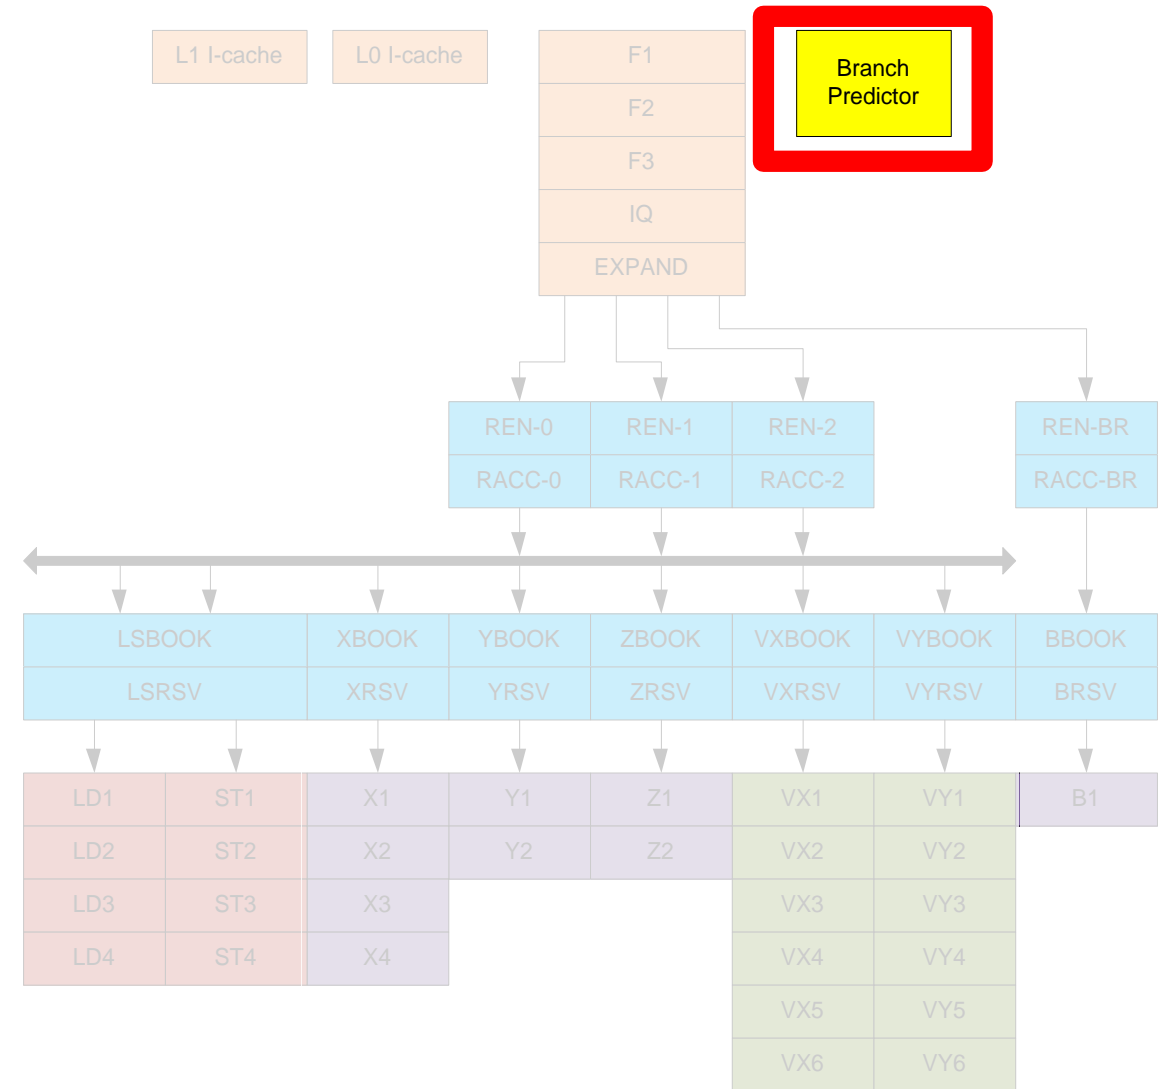
# Falkor CPU

- Heterogeneous pipeline providing optimal performance per unit power
  - Variable-length pipelines tuned per function
  - Minimizes idle hardware
- 4-issue
  - 3 instructions + 1 direct branch
- 8-dispatch



# Branch Prediction

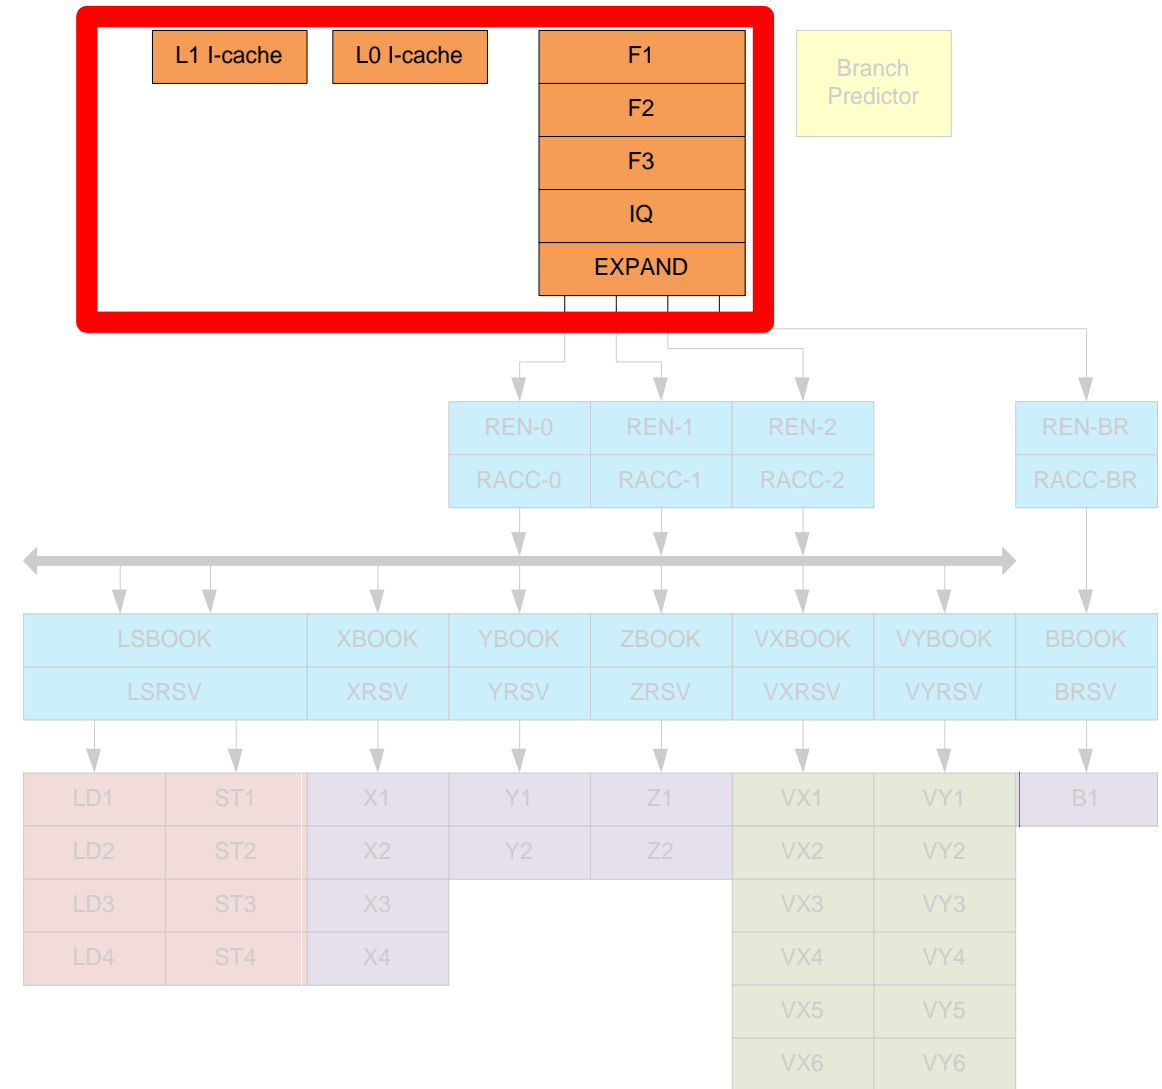
- 0-1 cycle penalty for almost all predicted taken branches
- 16-entry BTIC (branch target instruction cache)
  - Supports 0-cycle branch penalty
- Multi-level BTAC (branch target address cache) for indirect branches
  - 16-entry level-0 BTAC
  - 256-entry level-1 BTAC
  - PC-relative branches utilize I-cache as BTAC
- 16-entry link stack
- Multi-level BHT (branch history table)
  - Multi-faceted scheme involving staged predictors





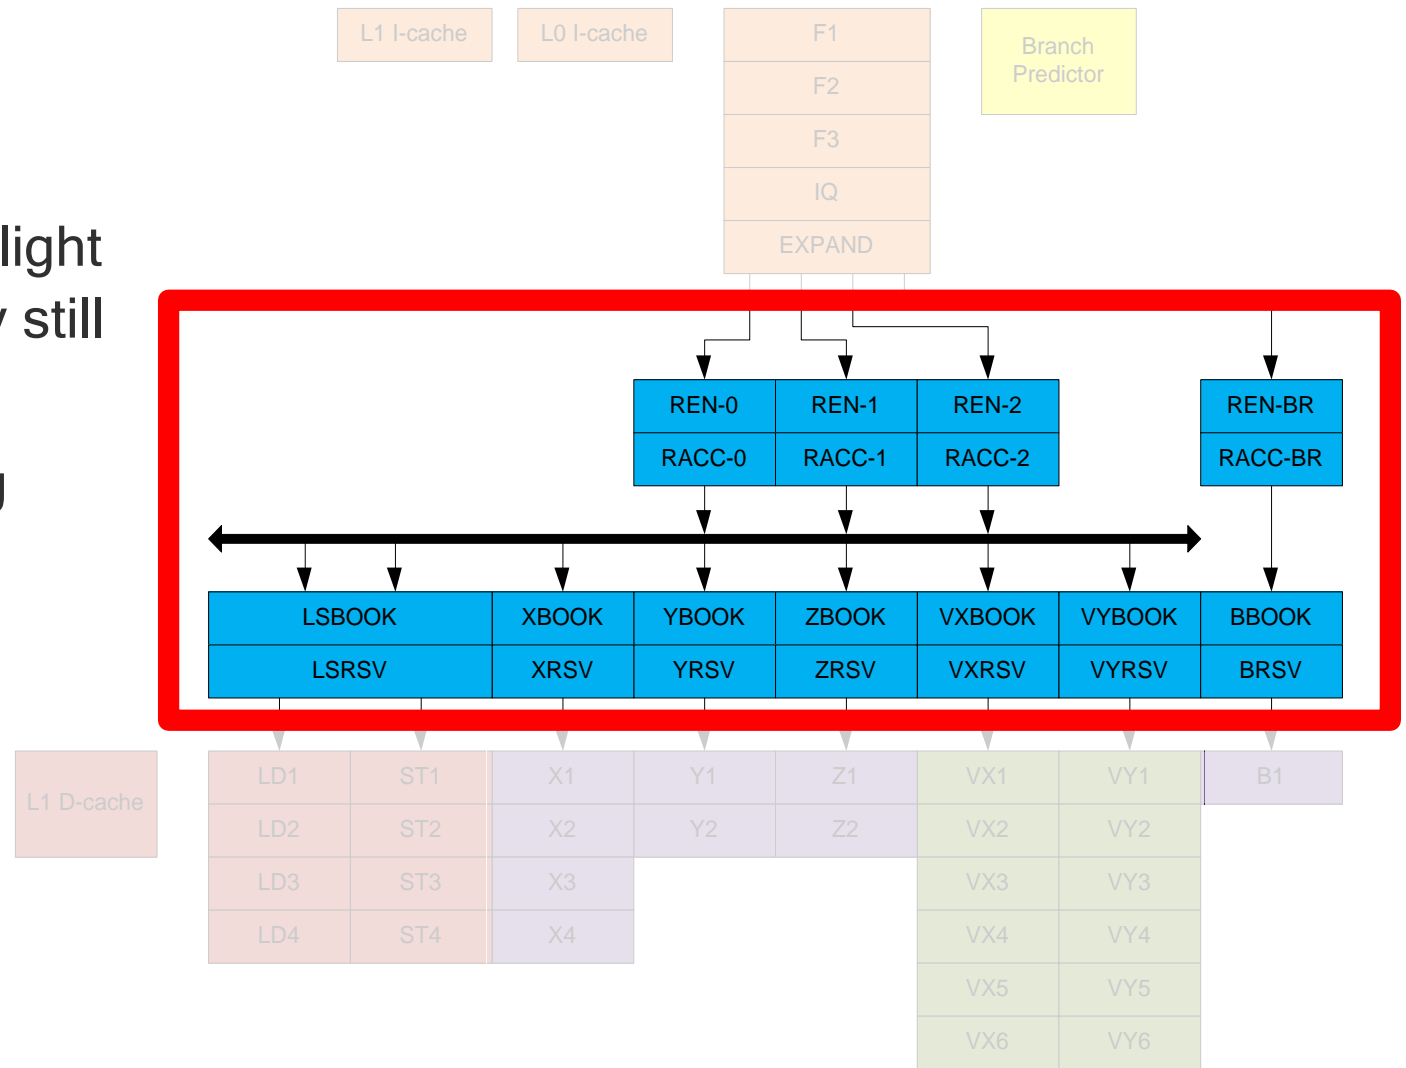
# Instruction Fetch

- Two-level I-cache topology
  - Key element in performance and performance/power efficiency advantage
  - L0 and L1 caches are exclusive
- L0 I-cache
  - 24KB, 64-byte lines, 3-way
  - Way-predicted
  - Parity with auto-correct
  - 0-cycle penalty for L0 hit
- L1 I-cache
  - 64KB, 64-byte lines, 8-way
  - Parity with auto-correct
  - 4-cycle penalty for L0 miss / L1 hit
  - Hardware prefetch on L1 miss
- Fetches up to 4 instructions per cycle
  - Fetch group can span cache lines
- Instructions are decoded and expanded into micro-ops
  - Most instructions map to a single micro-op



# Rename (REN), Register Access (RACC), and Reserve (RSV)

- 256-entry rename/completion buffer
- 76-instruction dispatch window
- Up to 128 uncommitted instructions in flight
  - Additional committed instructions may still be waiting on retirement
- Out-of-order dispatch of branches, ALU operations, loads, stores
- Up to 4 instructions retired per cycle

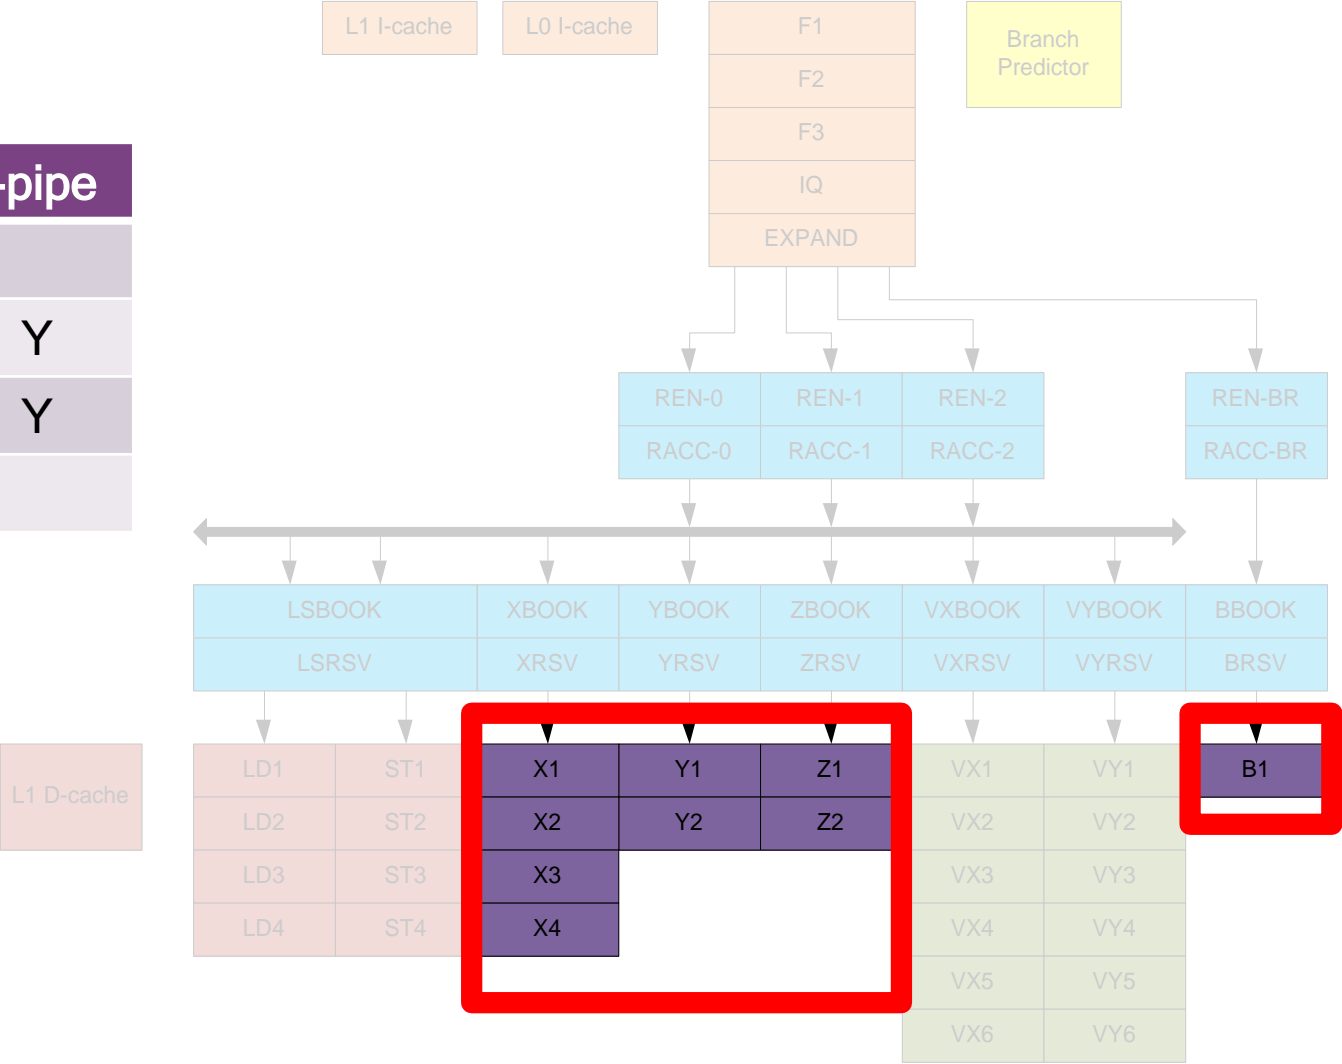


# Integer and Branch Execution

- Heterogeneous execution units for integer ALU operations and branches

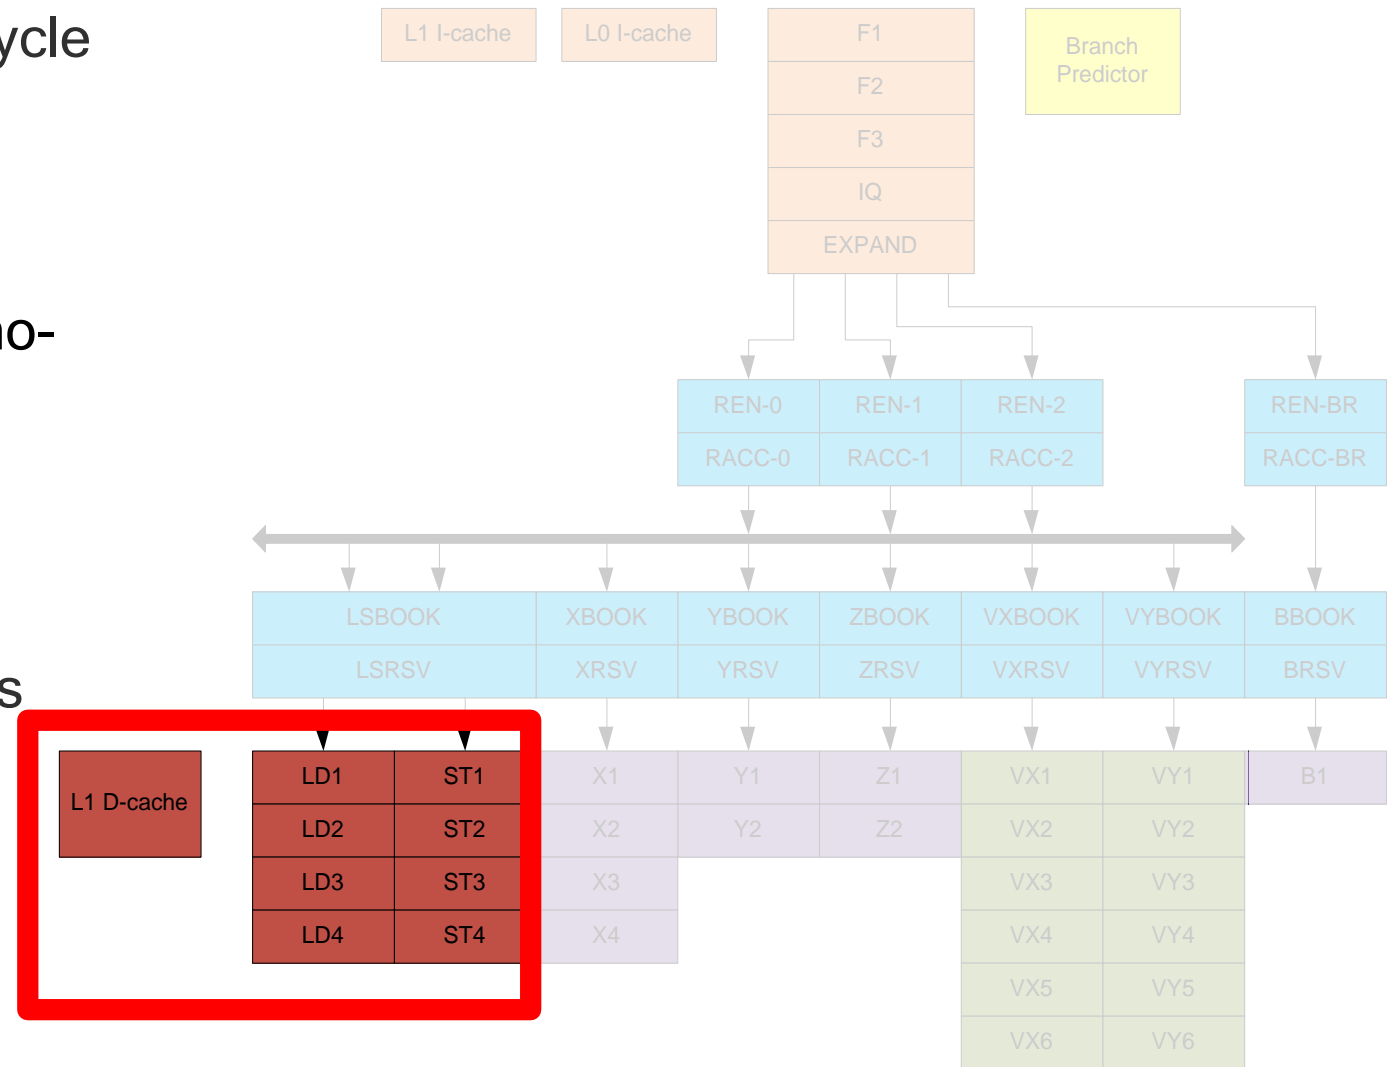
Operation	B-pipe	X-pipe	Y-pipe	Z-pipe
Direct branch	Y			
Indirect branch				Y
Simple ALU		Y	Y	Y
Multiplies		Y		

- Pipeline length sized based on operation



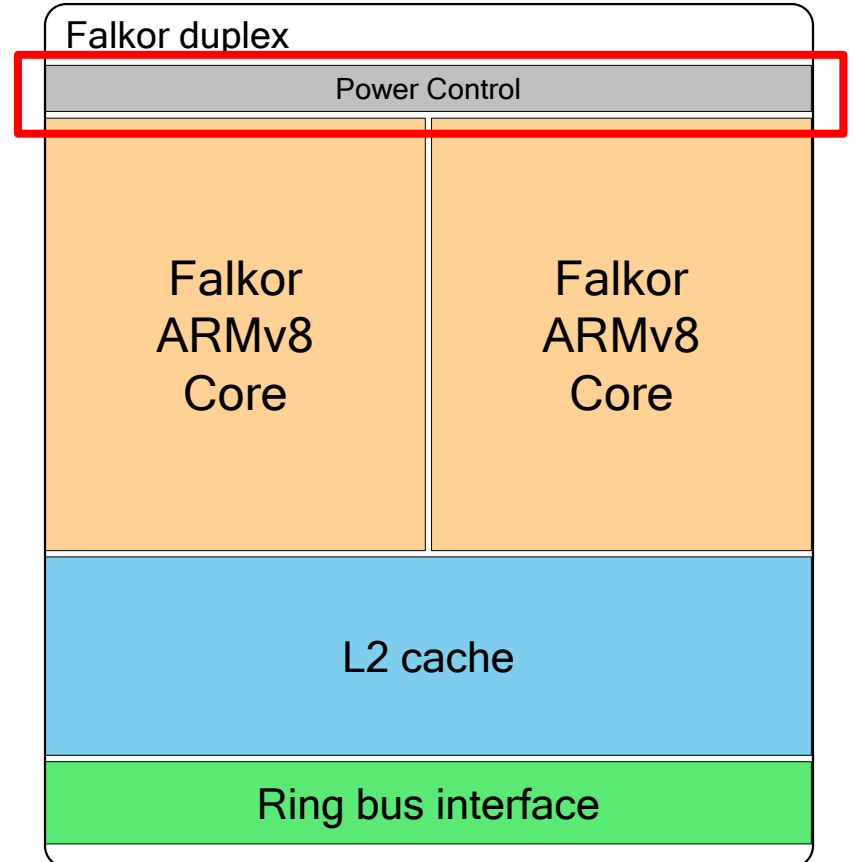
# Load/Store Execution

- 128 bits load and 128 bits store per cycle
- L1 data cache
  - 32KB, 64-byte lines, 8-way
  - 3-cycle latency for L1 hit
  - Write-through, read-allocate, write-no-allocate
  - Split virtual and physical tags
  - Parity with auto-correct
- Hardware data prefetch engine
  - Prefetches for L1, L2, and L3 caches
  - Detects stride patterns
- TLBs
  - 64-entry L1DTLB
  - 512-entry "final" L2TLB
  - 64-entry "non-final" L2TLB
  - 64-entry Stage-2 TLB



# Power Management

- Independent power states for each of CPUs and L2
- Each CPU is powered by a block head switch (BHS) or low-dropout regulator (LDO) from shared supply rail
  - Light sleep: gate off CPU clock
  - Voltage retention: registers and caches retain state
  - Register retention: register state retained using chip power rail
    - Caches and logic are switched off
  - Collapse: register and L1 cache state not retained
- L2 controller
  - Low-power states similar to CPU
  - L2 may auto-clock gate even when CPUs are active
  - L2 may enter retention or collapse state if both CPUs are in low-power states
- Entry/exit to/from low-power states controlled by hardware state machines
  - Minimizes entry/exit latency



# Qualcomm Centriq 2400 SoC Overview

## L3 Cache

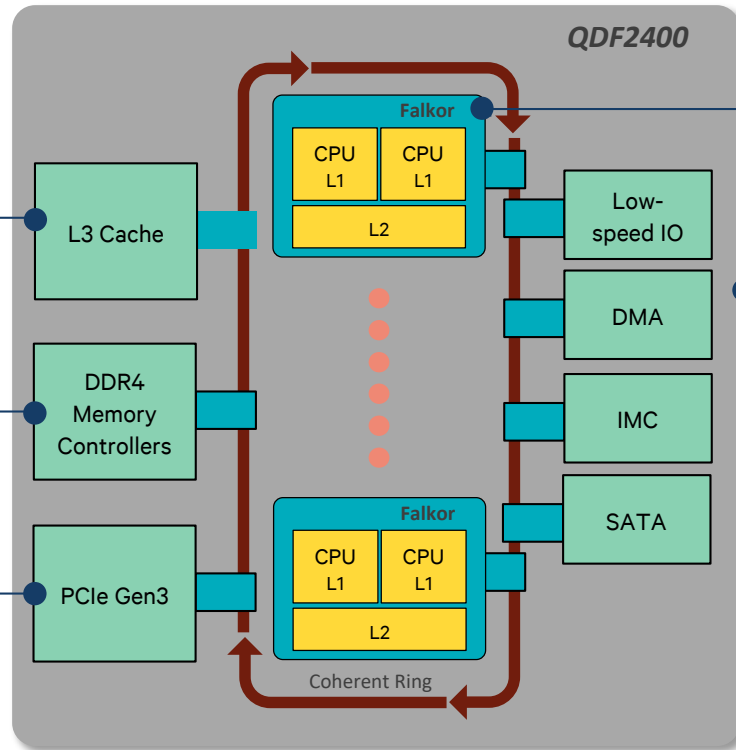
Large distributed unified L3 w/ECC

## DDR4 Memory

6 Channels w/ECC  
Bandwidth Compression  
2667 MT/s  
RDIMM, LRDIMM  
1 or 2 DIMMs per Channel

## PCIe Gen3

32 Lanes



## CPU Subsystem

Falkor cores based on ARMv8  
48 cores (24 duplexes)  
Unified L2 cache w/ECC

## SoC

Integrated "south bridge" features  
DMA, SATA, USB, I2C, UART, SPI, GPIO  
SBSA Level 3 Compliant

## Package

55mm x 55mm LGA  
Socketed

# L3 Quality of Service (QoS) Extensions

## *Shared Resource Contention Impacts QoS*


- **Distributed L3 Cache**
- **Limited/No Allocation Policy Enforcement for Data**

## QoS Extensions:

- Hardware Abstracted QoS Domain Identifier
  - Per Client (Core/Virtual Machine, IO/Virtual Function)
- Per-Resource Monitoring and Way-based Allocation
  - Monitor Utilization per QoSID per L3
  - Policy Enforcement per QoSID per L3
    - **Instruction/Data Granularity**
  - Fine-Tune Cache Allocation per Thread or Class of Threads

QoSID	Way Use Enable Mask (20-way)
0	0xFFFFF
1	0xF0000
2	0x0F000
...	
255	0xFF000

***Improved cache utilization and per-workload performance (lower application latency) for critical workloads.....***



# Memory Bandwidth Compression

## Constrained Memory Bandwidth

- **Channel limited peak Bandwidth**
- **Limited number of DDR Channels**

### Bandwidth Compression:

- Proprietary algorithm
- Inline compression w/in Memory Controllers
  - Fully transparent to software
- Compress 128B line to 64B when possible
- ECC is encoded with compression bit
- Very low latency decompression
  - 2 – 4 cycles
- Effective on compressible bandwidth intensive workloads
- Performance improvement varies with workload characteristics

Uncompressed Memory  
(128B Lines)

0a	0b	1a	1b
2a	2b	3a	3b
4a	4b	5a	5b
6a	6b	7a	7b
8a	8b	9a	9b
Aa	Ab	Ba	Bb



Compressed Memory

0		1	
2a	2b	3	
4		5a	5b
6a	6b	7a	7b
8		9a	9b
A		Ba	Bb

***Increased effective memory bandwidth  
and reduced power for compressible  
workloads.....***

### Memory Access Stream – w/ Bandwidth Compression

0	1	2a	2b	3	4	5a	5b	6a	6b	7a	7b	8	9a	9b	A	Ba	Bb
---	---	----	----	---	---	----	----	----	----	----	----	---	----	----	---	----	----

### Memory Access Stream – w/o Bandwidth Compression

0a	0b	1a	1b	2a	2b	3a	3b	4a	4b	5a	5b	6a	6b	7a	7b	8a	8b
----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----



# Secure Boot

---

- Immutable Boot ROM
  - Primary Boot Loader code resident in on-chip ROM
  - Contains code to authenticate external Firmware/Software
  - Establishes Root of Trust
- Security Controller / Fuse Block
  - Selection of public key
    - Qualcomm public key (from Boot ROM)
    - OEM public key
    - Customer public key (hash)
  - Authentication of secondary and tertiary Boot Loaders
- Integrated Management Controller
  - Dedicated processor for boot sequencing
  - Authenticates and anti-rollback checks Boot Loaders
  - Accelerates SHA portion of digital signature algorithm
    - Firmware performs RSA public key operations

# Summary & Status

---

- **Qualcomm Centriq™ 2400 Processor is the industry's first 10 nm server CPU**
- 5th-generation custom core design
  - Specifically optimized for server applications
- ARMv8-compliant AArch64 only
- Targeting leading-edge Performance with Performance per Watt leadership
- Motherboard specification submitted to Open Compute Project
  - based on the latest version of Microsoft's Project Olympus
- Running Windows Server and multiple versions of Linux
- Chip is being sampled at multiple datacenters
- On track for production by end of 2017



# Thank you

---

Follow us on:    

For more information, visit us at:

[www.qualcomm.com](http://www.qualcomm.com) & [www.qualcomm.com/blog](http://www.qualcomm.com/blog)

Nothing in these materials is an offer to sell any of the components or devices referenced herein.

©2017 Qualcomm Technologies, Inc. and/or its affiliated companies. All Rights Reserved.

Qualcomm is a trademark of Qualcomm Incorporated, registered in the United States and other countries, Qualcomm Centriq and Falkor are trademarks of Qualcomm Incorporated. Other products and brand names may be trademarks or registered trademarks of their respective owners.

References in this presentation to “Qualcomm” may mean Qualcomm Incorporated, Qualcomm Technologies, Inc., and/or other subsidiaries or business units within the Qualcomm corporate structure, as applicable.

Qualcomm Incorporated includes Qualcomm’s licensing business, QTL, and the vast majority of its patent portfolio. Qualcomm Technologies, Inc., a wholly-owned subsidiary of Qualcomm Incorporated, operates, along with its subsidiaries, substantially all of Qualcomm’s engineering, research and development functions, and substantially all of its product and services businesses, including its semiconductor business, QCT.



# Glossary

---

- SoC - System-on-Chip
- SBSA - Server Base System Architecture
- LGA - Line Grid Array
- SATA - Serial Advanced Technology Attachment
- USB - Universal Serial Bus
- I2C - Inter-Integrated Circuit
- UART - Universal Asynchronous Receiver/Transmitter
- SPI - Shared Peripheral Interrupt
- GPIO - General Purpose Input Output
- RDIMM - Registered (Buffered) Dual Inline Memory Module
- LRDIMM - Load Reduced Dual Inline Memory Module
- DDR - Double Data Rate